A NON-BLOCKING DEVICE FOR SWITCHING ATM CELLS

The field of the invention is that of digital data transmission. To be more precise, the invention concerns switching ATM (Asynchronous Transfer Mode) cells between a plurality of incoming channels and a plurality of outgoing channels, in particular in the case of variable bit rate ATM services.

The general principle of ATM transmission is well known and many techniques for implementing switching nodes between inputs and outputs have already been proposed. One criterion for estimating the efficacy of such techniques is cell loss rate, in other words the risk of blocking within a switching node.

This risk is relatively high in prior art techniques, in particular in the context of variable bit rate services. It can happen that cells accepted at an input of the switching node cannot be transferred to the required output because of excessive congestion on one of the internal links of the switching node.

The techniques most widely used at present are based on the Clos structure which is described for example in "A study of Non-Blocking Switching Networks" by C.Clos (The Bell System Technical Journal, pp 406-424, March 1953).

Figure 1 shows a Clos network. It has three stages:

- an inlet stage 11 comprising a plurality of matrices 111 with R inputs 112 and K outputs 113;
- a central (intermediate) stage 12 comprising K matrices 121 each connected to one of the outputs 113 of each of the input matrices 111; and
- an outlet stage 13 comprising the same number of matrices 131 as the inlet stage 11, each of the matrices 131 having K inputs 132 respectively connected to the K matrices 121 of the central stage 12 and R outputs 133.

Using a plurality of stages minimizes both the number of connection points and the size of the matrices.

The above technique is effective if the switched

calls are at fixed bit rates. It can then be shown that it is sufficient for K to be greater than or equal to 2R-1 for the network to be non-blocking. In a non-blocking network any free input (i.e. any input able to accept a new call) can be connected to any free output.

The essential problem with the ATM technique is that it is called upon to transmit calls with variable bit rates which can sometimes require a very low or even zero bit rate and which can at other times require a high bit rate up to some given peak bit rate.

When a Clos structure is used, this leads to considering each call to have a fixed bit rate equal to its peak bit rate. Clearly the above approach is particularly ineffective, especially when calls come in bursts, and feature long time periods in which the bit rate is low. A high bandwidth is then reserved permanently, to no good purpose.

Consider for example two calls with low levels of activity but requiring high bandwidths when they are active. If there is a low probability that their times of activity are the same they can use the same path. On the other hand, on the basis of the peak bit rate, cohabitation on the same path is impossible.

For a given switching structure it would be possible at all times to test the validity of the selected paths against the specific multiplexing law employed. However, there is no a priori method of quantifying the blocking of a network of the above kind, given the infinite number of possible combinations of traffic on the incoming and outgoing links.

One object of the invention is to alleviate the drawbacks of the prior art.

To be more precise, a first object of the invention is to provide a non-blocking device for switching ATM cells. In other words, the object of the invention is to provide a device of the above kind that can provide switching between any free input and any free output.

The invention applies in particular to switching variable bit rate calls multiplexed at the input and/or at the output. For this type of system a traffic acceptance law is defined guaranteeing a probability of saturation below a given threshold (conventionally 10^{-x}). A call can then be accepted on a link if the total traffic has an acceptable saturation probability.

Given this approach, the object of the invention is to provide a switching device that guarantees a saturation probability on the intermediate links that is less than or equal to that accepted (in accordance with the same law) on the input and output links.

In other words, an object of the invention is to provide a device of the above kind that is statistically non-blocking.

The above objects, together with others that become apparent below, are achieved according to the invention with the aid of a device for switching ATM cells establishing a single path per virtual circuit, having N.R inputs and N.R outputs, N and R being two integers not less than two, the device comprising at least two stages, including an inlet stage having R.N sets of Q outputs and an outlet stage having R.N sets of Q′ inputs,

characterized in that for the flow of data carried by any intermediate link that is part of the single path set up between an input and an output to be a subset of the incoming flux at that input and also a subset of the outgoing flux at that output, each input of the inlet stage can be connected to an output of the inlet stage which can be selected only from Q outputs exclusively associated with that input; and

in that each output of the outlet stage can be connected to an input of the outlet stage which can be selected only from Q′ inputs of the outlet stage exclusively associated with that output.

Accordingly, the device in accordance with the invention uses a tree structure with perfect meshing at

link level.

The same multiplexing rule is advantageously used on the incoming links and on the outgoing links of all said stages.

It can easily be shown that in this case each intermediate link can convey only a subset (or in an extreme situation the same set) of traffic accepted on each incoming link and each outgoing link. It is therefore certain that the saturation probability is less than or equal to that of the input and output links.

In other words, the switching network is "transparent" in terms of blocking probability. If traffic (or a service) can be accepted at an input and an output, it is certain that it can be transmitted from one to the other.

The invention is independent of the multiplexing law adopted at the inputs and the outputs.

Various architectures can be considered for a device of the above kind. For example it may comprise only an inlet stage and an outlet stage each comprising N switching matrices, and be characterized:

in that, Q being equal to N, each matrix of the inlet stage has R inputs and R.N outputs organized into R sets of N outputs, each set corresponding to a respective one of the R inputs; in that each input of that matrix can be connected to an output of that matrix which can be selected only from N outputs of the set of outputs corresponding to that input;

in that, Q' being equal to N, each matrix of the outlet stage has R outputs and N.R inputs; and in that each output of that matrix can be connected to an input of that matrix which can be selected only from R.N inputs of that matrix; and

in that each of the N outputs of each set of outputs of the first stage is connected to an input of a respective one of the N matrices of the outlet stage.

The above structure nevertheless pre-supposes large

matrices for the second stage. Another embodiment comprises an inlet stage, a central stage, and an outlet stage and is characterized:

- in that, Q being equal to R, the inlet stage comprises N matrices each having R inputs and R^2 outputs, those outputs being organized into R sets of R outputs each corresponding to one of said R inputs, and in that each input of that matrix can be connected to an output of that matrix which can be selected only from R outputs of the set of outputs corresponding to that input;
- in that the central stage comprises R sets of R matrices each having N inputs and N outputs, the R outputs of each set of outputs of the inlet stage being connected to inputs belonging to the same set of R matrices of the central stage; and
- in that, Q′ being equal to R, said outlet stage comprises N matrices each having R^2 inputs and R outputs, those R^2 inputs being organized into R sets of R inputs, each set respectively corresponding to one of those R outputs; and in that each output of that matrix can be connected to an input of that matrix which can be selected only from R inputs of the set of inputs corresponding to that output; and in that the R inputs of each set are respectively connected to R outputs respectively belonging to the R sets of matrices of the central stage.

It is therefore possible to use smaller, conventional matrices.

In another embodiment the device of the invention can comprise an inlet stage, a central stage, and an outlet stage and be characterized:

- in that Q and Q′ are equal to R,
- in that the central stage includes R^2 matrices,
- in that the inlet stage and the outlet stage each comprise R.N switching matrices,
- in that the matrices of the inlet stage and the matrices of the central stage are organized into R sets

each including N matrices of the inlet stage and R matrices of the central stage and the matrices of the outlet stage are organized into N sets of R matrices;

- in that each of the R.N matrices of the inlet stage has a single input and R outputs,

- in that each of the R^2 matrices of the central stage has N inputs and N outputs, the N inputs being respectively connected to an output of each of the matrices of the inlet stage that belong to the same set of matrices; and

- in that each of the R.N matrices of the outlet stage has R inputs and a single output, those R inputs being connected to outputs respectively belonging to the R sets of matrices of the central stage and of the inlet stage.

A new type of matrix specifically adapted to the invention is then used.

N and R are then preferably chosen so that $N = 2.R^2$ if the device has three stages. This structure is the most effective one, in particular for assuring homogeneity between narrowband networks and broadband networks, since this allows the same matrices to be used.

Other features and advantages of the invention become apparent on reading the following description of preferred embodiments of the invention given by way of illustrative and non-limiting example only and from the accompanying drawings, in which:

- Figure 1, already discussed in the preamble, shows the known principle of a Clos network;

- Figure 2 is a block diagram of a two-stage switching network in accordance with the invention;

- Figure 3 shows a different embodiment of a switching network with three stages;

- Figures 4A and 4B show an advantageous structure based on the network from Figure 3 allowing broadband transfers (Figure 4B) and narrowband transfers (Figure 4A); and

Sub B1

- Figure 5 also shows a three-stage switching network in accordance with the invention, including 1×R matrices.

The invention concerns a non-blocking ATM cell switching device or network using statistical multiplexing, in particular for variable bit rate ATM services.

Figure 2 shows the general principle of the invention in the case of a network having two stages, an inlet stage 21 and an outlet stage 22. The inlet stage 21 comprises N input matrices 211_1 through 211_N which each receive R incoming (or input) links 212_1 through 212_R and which each have R.N intermediate links $213_{1,1}$ through $213_{R,N}$. The outlet stage 22 comprises N output matrices 221_1 through 221_N which each also receive R.N intermediate links $222_{1,1}$ through $222_{N,R}$ and deliver R outputs 223_1 through 223_R.

Each link 212_i, 213_i and 223_k conforms to the same traffic acceptance law (for example a saturation probability less than 10^{-x}).

Sub B2

The input matrices 21_i are organized so that the incoming flux of data at each input 212_i can be directed to any matrix 22_i of the outlet stage. In other words, a tree structure is used which can define N possible connections for each input and no more than N (N is the number of matrices in the outlet stage in this embodiment).

Similarly, each matrix 221_i of the outlet stage can receive data from each input 212_i of the inlet stage. The flux conveyed by each of the intermediate links 222_k reaching a matrix 221_i can be transmitted to any of its outputs 223_i.

There is no risk of blocking on the intermediate links 222_k. They always carry:

- a subset (or possibly the complete set) of traffic accepted on the input link from which they originate; and

- a subset (or possibly the complete set) of traffic accepted on the output link to which they lead.

By definition the traffic accepted on these input and output links conforms to a predefined multiplexing law.

Note however that the Figure 2 structure requires large matrices in the second stage because of the large number of possible connections.

It is possible to implement the invention using smaller matrices by constructing a three-stage switching network like that shown in Figure 3.

The inlet stage 31 comprises N matrices 311_1 through 311_N each associating R inputs 312_i through 312_R with R^2 outputs $313_{1,1}$ through $313_{R,R}$. Symmetrically, the outlet stage 33 also has N matrices 331_1 through 331_N which each receive R^2 inputs $332_{1,1}$ through $332_{R,R}$ and each deliver R outputs 331_1 through 333_R. The intermediate (central) stage 32 comprises R groups of R N×N matrices $321_{1,1}$ through $321_{R,R}$.

Sub B3

The links between the various stages are organized so that the flow received by each input of a matrix 311_i of the inlet stage can be transmitted to any of the corresponding R matrices $321_{i,1}$ through $321_{i,R}$. Similarly, each output matrix 311_i can receive data from each matrix 321_k of the central stage. To be more precise, each output link 331_{ii} can receive data from any of the R matrices $321_{1,i}$ through $321_{R,i}$.

As in Figure 2, and for the same reasons, the Figure 3 network is non-blocking.

For example, the above network can be implemented using conventional 16×16 matrices. For a network with 16×3 = 58 622 Mb/s ATM links, equivalent to a prior art network with 58 links (or 16×4 limited to 80%), the following are used:

- 16 3×9 input/output matrices, and
- 9 16×16 central matrices,

i.e. 25 16×16 matrices rather than the 24 in the

prior art network. The device in accordance with the invention therefore requires one additional matrix. On the other hand, it can multiplex all types of service (including services at 600 Mb/s). For services with a peak bit rate of 10 or 20 Mb/s, for example, such multiplexing allows up to five times more traffic to be connected.

Accordingly the device in accordance with the invention makes it possible to envisage a bit rate of 3000 Mb/s whereas a network operating on the basis of the peak bit rate alone, without allowing for the "bursty" nature of the traffic, can of course switch only 600 Mb/s per link. In practice the link is further restricted to a maximum bit rate of 150 Mb/s (otherwise the expansion would have to be even greater).

The values N and R are advantageously chosen so that $N = 2R^2$ if the device has three stages. This provides a simple way of assuring homogeneity between the "narrowband" network and the "broadband" network.

This is shown in Figures 4A and 4B. In this example N = 32 and R = 4. The matrices 41 of the inlet stage and the matrices 42 of the outlet stage are 16×16 matrices. The matrices 43 of the central stage are 32×32 matrices. Figure 4A shows a "narrowband" network switching device. Each of the 16 inputs of a matrix of the first stage 41 can be connected to any of the 16 outputs of that matrix. Each of the 16 outputs of a matrix of the last stage 42 can be connected to any of the 16 inputs of that matrix.

Figure 4B shows a "broadband" network device, of the type shown generally in Figure 3, for which N = 32 and R = 4. Only four (from sixteen) inputs of each matrix of the first stage 41 are used. Each can be connected to one of four outputs of that matrix exclusively associated with that input. Only four (from sixteen) outputs of each matrix of the last stage 42 are used. Each can be connected to one of four inputs of

that matrix exclusively associated with that output.

It is a simple matter to derive a "narrowband" network from a "broadband" network, as shown in Figure 4A, because all that is required is to use 16 inputs per matrix 41 and 16 outputs per matrix 42 for the "narrowband" network or only 4 inputs per matrix 41 and 4 outputs per matrix 42 for the "broadband" network. The switching device as a whole remains unchanged.

In practice the "broadband" network (Figure 4B) is derived from the "narrowband" network (Figure 4A). The matrices are then used incompletely (only 4 inputs or 4 outputs) and restricted routing is employed in the input and output matrices.

Figure 5 shows another implementation of the technique of the invention. Instead of using prior art (N×N) matrices incompletely, specially designed 1×R matrices are used (associating an input with R outputs or, symmetrically, R inputs with one output).

Using this technique there are R input blocks 41_1 through 41_R and N output blocks 42_1 through 42_N.

Each input block 41_i comprises:

- N input matrices 411_1 through 411_N each having a single input 412_1 through 412_N and R outputs $413_{1,1}$ through $413_{N,R}$. In this embodiment each input is associated with R links respectively corresponding to the R matrices of the next stages; and

- R central matrices 414_1 through 414_R each having N inputs and N outputs. Each input corresponds to a distinct input matrix.

Each of the N output blocks 42_i comprises R matrices 421_1 421_R each having a single output 422_1 422_R and receiving R links $423_{i,1}$ through $423_{i,R}$ from the respective R input blocks 41_1 41_R.

In this structure each link 413_i carries only a subset of traffic from the associated input and each link 423_i carries only a subset of traffic from the associated output. This structure highlights how the invention

functions, namely by demultiplexing input traffic and then demultiplexing outgoing traffic so that certain traffic mixes are avoided.

Note that the matrices used are of optimum size.

Figure 5 also shows clearly that the structure of the device of the invention very readily enables broadcasting, which is particularly important in the case of "broadband" services.

Of course, the use of 1×R matrices described here in a three-stage system can readily be transposed to the two-stage device shown in Figure 2.